NaSPA technical SUPPORT

Supporting Enterprise Networks and Operating Environments

APRIL 1996
VOLUME 4, NUMBER 4

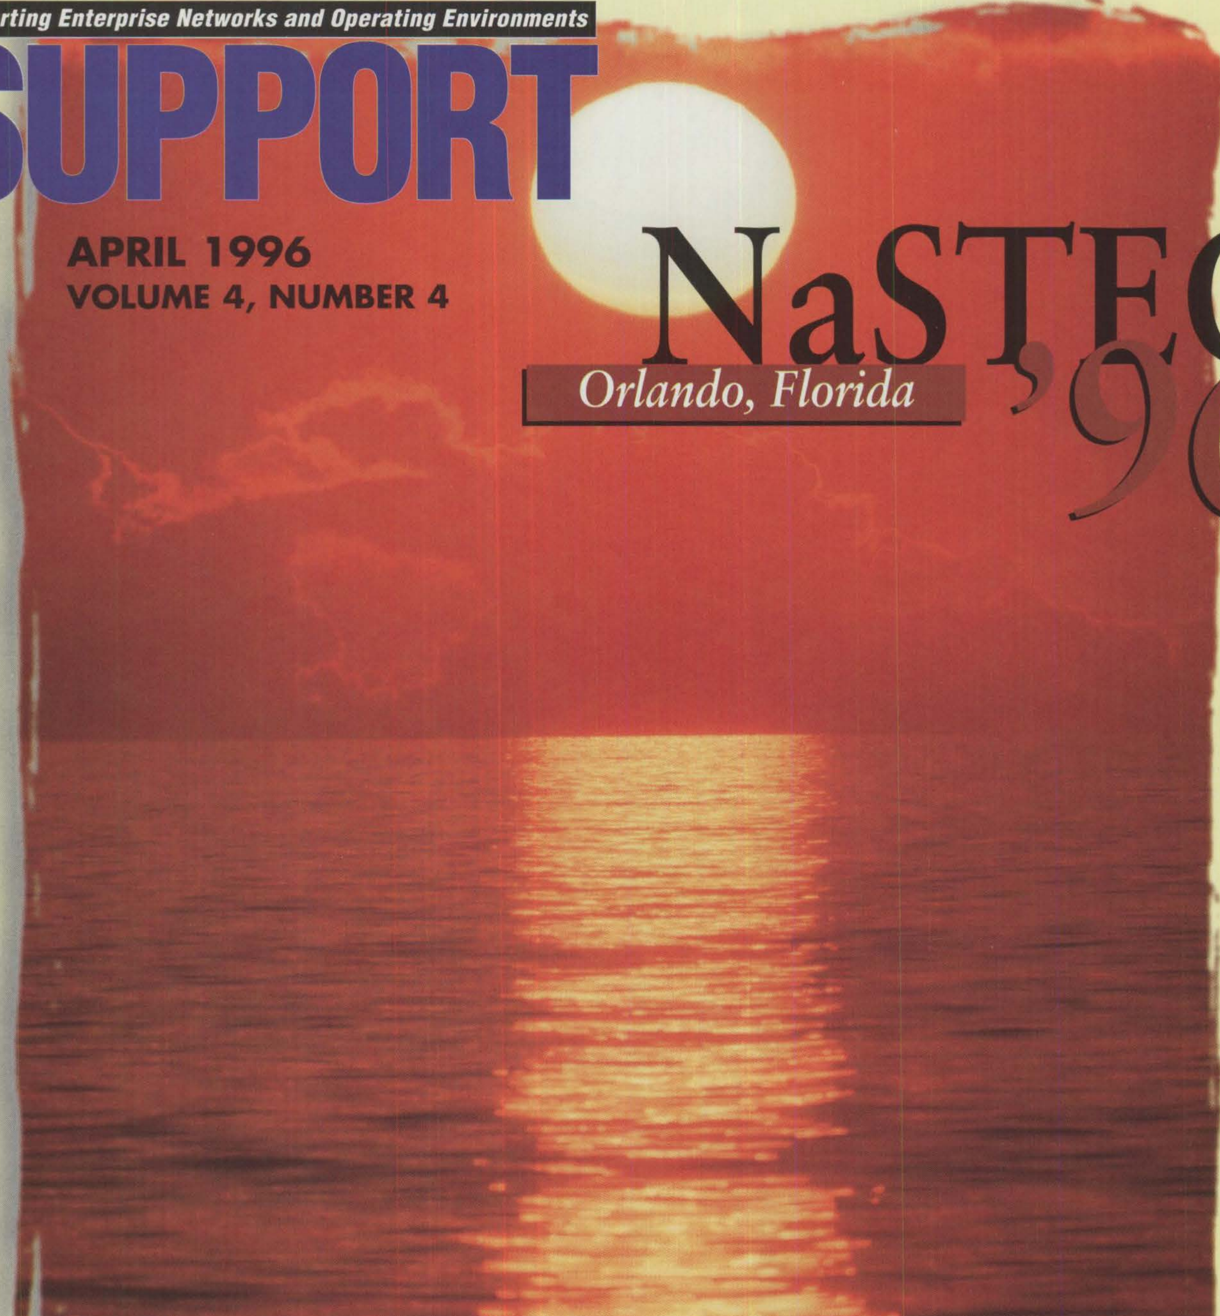

NaSTEC '96

Orlando, Florida

http://www.nascom.com/

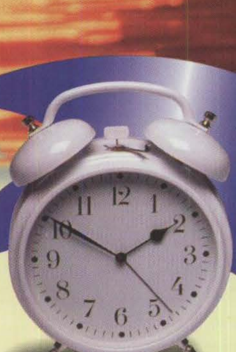

HURRY!

Time is Running Out for NaSTEC '96.

See page 40 for details.

technical® SUPPORT
Supporting Enterprise Networks and Operating Environments

APRIL 1996
VOLUME 4, NUMBER 4

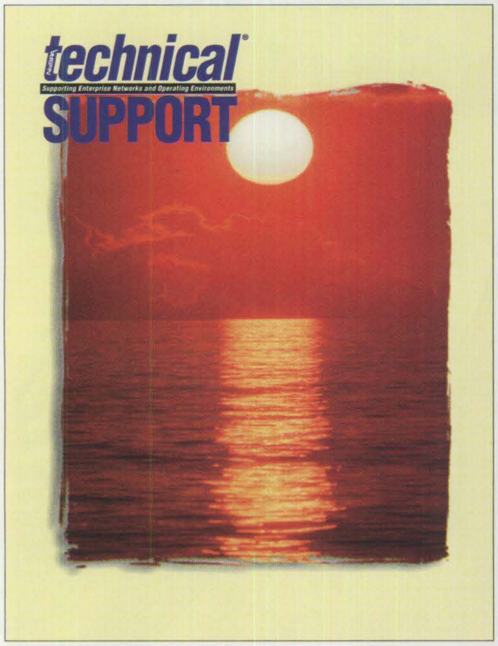

FEATURES

SYSTEM STRATEGIES

NaSPA Mission Statement:

The mission of NaSPA, Inc., a not-for-profit organization, shall be to serve as the means to enhance the status and promote the advancement of all corporate computing technical professionals; nurture member's technical and managerial knowledge and skills; improve member's professional careers through the sharing and dispersing of technical information; promote the profession as a whole; further the understanding of the profession and foster understanding and respect for individuals within it; develop and improve educational standards; and assist in the continuing development of ethical standards for practitioners in the industry.

NaSPA, Inc. is a not-for-profit, independent corporation and is not owned in whole or in part by any manufacturer of software or hardware. All corporate computing professionals are welcome to join NaSPA, Inc. Membership rates are $29.95/year (USA), $44.95/year (Canada) and $54.95/year (all other countries). $19.98 of your annual dues is allocated to the publication ***NaSPA Technical Support*** and is non-deductible therefrom.

NaSPA Technical Support (ISSN 1079-3135) (IPM Agreement Number 0806773) is published monthly by Technical Enterprises Inc., 7044 S. 13th Street, Oak Creek, WI 53154-1429. Second-class postage paid at Oak Creek, WI and additional mailing office. POSTMASTER: Send address changes to ***NaSPA Technical Support***, 7044 S. 13th Street, Oak Creek, WI 53154-1429.

INTEROPERABILITY INSIGHTS

ATTENTION READERS!

For faster response to your advertising and information inquiries, we have moved the Reader Response Card inside the magazine to page 69. This will allow us to process your information quicker and provide you with the requested materials sooner.

COLUMNS

DEPARTMENTS

Cutting the FAT

BY MICHAEL NORTON

One of the consequences of IBM's misguided strategy to compete with Microsoft in the desktop market was the introduction of OS/2 systems on numerous machines using the antiquated File Allocation Table (FAT) system. The new user's introduction to OS/2 was, unfortunately, a blue screen proffering an "easy" install option (the default) over the "advanced" install option. This included an ominous warning about needing to be a technical dweeb/2 to even attempt such a thing (with the implication, don't bother calling IBM if you encounter any difficulties because you were warned). This meant that most users installed ATDT942-7278 OS/2 with the dual–boot option in their existing DOS/Windows partition, which, by definition, used the FAT file system.

OS/2'S RELIANCE ON EXTENDED ATTRIBUTES

OS/2 should be run on an HPFS partition. The OS/2 operating system relies heavily on the use of extended attributes (EAs) to maintain its object–oriented associations. Your desktop, for example, consists largely of extended attributes contained in the DESKTOP directory structure. Extended attributes are "built in" to the HPFS file system; for FAT they are merely an afterthought.

Let's look at how FAT handles extended attributes. If you add an extended attribute to a file (i.e., the .LONGNAME standard extended attribute), the actual data is not written to (or even with) the file. Instead, it is written to a free allocation unit on the volume as independent data. This data is then associated with the file by a pointer in an area of the file's directory entry that is unused by DOS. The allocation unit in the FAT is then assigned to the system file EA_DATA._SF to prevent the data from being used by the operating system for new or appended files. Note that an entire allocation unit is assigned to the EA. Thus, simply assigning a long file name to a file consumes 8 KB of disk space on volumes using 8 KB clusters (which is fairly typical)! This is the reason many OS/2 users running on the FAT file system report absurdly large EA_DATA._SF files.

There are, indeed, a number of other advantages to running HPFS instead of the FAT file system, but, alas, print space, like disk space, is finite and my soap box is only so tall. If you are using the FAT file system, I hope you will reconsider, especially on your boot volume. You will experience fewer problems with HPFS.

EAS ON AN HPFS VOLUME

Contrast this cumbersome and inefficient method of handling EAs to that employed by HPFS. In HPFS, each file is assigned its own header, a one–sector construct called the F-node, which immediately precedes the file data on the volume. The F-node contains the file name (including the long name) and allocation information in addition to the file's extended attributes. Because of this intimate association between a file and its F-node, it is much more difficult to lose or corrupt EAs on an HPFS volume than on a FAT volume. Additionally, since typically only the 512–byte F-node is used to contain EAs (although more space may be allotted if required), there is a marked improvement in storage efficiency.

This improvement is not limited to EAs; the FAT file system is notoriously wasteful of disk space because of short–sightedness in its design. Any file system performs two vital functions: maintaining an "index" of each file's location on the volume, and managing free space to prevent existing data from being overwritten. Although it must have seemed ingeniously clever at the time, the design strategy taken by FAT to combine these two functions in a single allocation table has yoked FAT ever since. On FAT16 systems (there were earlier incarnations that were even worse), two bytes are assigned to each allocation unit — originally a disk sector. Each allocation unit points to the next allocation unit assigned to the file or indicates the current allocation unit is the last assigned to the file. The two–byte design implies that the volume size is limited to 32 MB, since the last allocation unit that could theoretically be pointed to is 0xFFFF. To break the 32 MB barrier, FAT increased the size of the allocation unit from the 512–byte sector to the cluster, which is typically comprised of 16 sectors. Thus, a one–byte file occupies more than 8 KB of disk space.

HPFS was one of those rare moments in the computer industry when mistakes were admitted and learned from. HPFS abandoned the strategy of using the allocation table to maintain file location information and manage free space.

Instead, HPFS uses a bitmap, with one bit representing each sector, to indicate sectors in use, and a file's allocation information is maintained in its F-node. This design allows HPFS to utilize the sector rather than the cluster as the basic allocation unit. The same one–byte file which consumes 8 KB under FAT uses 1 KB under HPFS (one 512–byte sector for the data and another sector for the F-node).

The decentralization of allocation information has another, more important consequence: decreased susceptibility to catastrophic data loss. The FAT is obviously the most heavily used area on a FAT volume, and is written and rewritten constantly. With each access, the odds increase of something going awry, and, when it does, the damage is not limited to the file currently being allocated, but potentially all files on the volume. Moreover, if the FAT is corrupted, recovery is only possible through tedious and prohibitively costly measures. With HPFS, the interspersing of allocation information virtually prevents such

(continued from page 66)

massive data loss, and allows the possibility of recovery of much, if not all, of the volume data through the use of disk utilities such as SofTouch System's GammaTech Utilities or Warp Speed's Graham Utilities.

There are, indeed, a number of other advantages to running HPFS instead of the FAT file system, but, alas, print space, like disk space, is finite and my soap box is only so tall. If you are using the FAT file system, I hope you will reconsider, especially on your boot volume. You will experience fewer problems with HPFS. ts

Was this column of value to you? If so, please circle Reader Response Card No. 47.

Michael Norton is the workstation division manager at SofTouch Systems, which provides both mainframe and PC software solutions. He has written mainframe manuals in addition to articles in a number of publications. You can reach him at mnorton@softouch.com.

(continued from page 65)

about this, it makes sense. The two MSEngines always arrive at the same results. If a particular event causes the primary MSEngine to crash, that same event will lead to the crash of the secondary MSEngine.

SFT III must be viewed as an added layer of protection against hardware errors. Many of us use mirror NetWare disk partitions to protect against hardware errors on our disk drives. SFT III takes this concept one step further by allowing us to mirror the entire server. ts

Was this column of value to you? If so, please circle Reader Response Card No. 46.

NaSPA member John E. Johnston is manager of technical support and communications for a major hospital in Pennsylvania. He designs and maintains cross- platform local and wide area networks utilizing NetWare, OS/2, DOS, and Windows.

NaSPA Manufacturer's Support Numbers

3COM	800-876-3266
3M	800-328-1360
Acculogic	714-454-2441
Adaptec	408-945-2550
Adobe	800-628-2320
Altsys (Macromedia)	415-252-9080
Apple	800-767-2775
Ares	415-578-9090
Artisoft	520-670-7000
Asante	800-622-7464
AST Research	800-727-1278
ATI Technologies	905-882-2626
Attachmate	800-688-3270
Berkeley	510-549-2300
Boca	407-241-8088
Borland	800-841-8180
Caere	408-395-8319
Calcomp	800-225-2667
Canon	800-423-2366
Central Point Software (Symantec)	503-690-8080
Claris	408-727-9054
(CMS)	
AmeriQuest Technologies	800-237-2707
Colorado	970-635-1501
Compaq	800-652-6672
Computer Associates	408-432-1764
Conner Peripherals	800-426-6637
Corel	800-818-1848
CPI ACC (Tech)	714-454-2441
Creative Labs	405-742-6622
Cyrix	800-462-9749
Dayna	800-531-0600
Daystar	770-967-2077
Delrina	416-441-1026
Deneba	800-622-6827
Diamond Multimedia	408-325-7100
Dlink	714-455-1688
(E-Machines) Radius	408-541-5700
Epson	800-922-8911
Farallon	510-814-5030
Fujitsu America	408-894-3950
FWB	415-833-4578
Global Village	408-523-1050
Goldstar	800-777-1192
Hayes	770-441-1617
Hercules	510-623-6050
Hewlett Packard	208-323-2551
Houston Instruments Summagraphics	800-444-3425
IBM Computers	800-426-7763
IBM/Lexmark	606-232-3000
Intel	503-629-7000
Iomega	800-456-5522
Irwin (Connor)	800-526-6637/800-426-6637
Kensington	800-535-4242
Keytronic	800-262-6006
Laser Tools	510-420-1319
Logitech	510-795-8100
Lotus Development	800-223-0303
Lotus Word Processing	800-421-6264
Lotus Request	800/346-3508
Luto America	800-445-8782
Macromedia	415-252-9080
MAG Innovision	800-827-3998
Magnavox & Philips	615-475-8869
Maxis	510-927-3905
Maxtor	800-262-9867
Maxtor hard drive 24-hr FAX back service	303-678-2618
Megahertz	801-320-7777
Microdyne	800-578-3270
Micronet	714-453-6060
Micropolis hard drive 24-hr FAX back service	818-709-3325
Microsoft	206-454-2030
Microsoft Windows 95 Support	206-635-7000
Microtek	310-297-5100
Mitsubishi	800-344-6352
Mountain	800-458-0300
Mountaingate	800-447-8302
Mouse Systems	510-656-1117
Nanao	800-800-5202
NEC Technologies	800-388-8888
Nikon	516-547-4311
Novell	800-638-9273
Okidata	800-654-3282
Ontrack	612-937-2121
Orchid Technology	510-683-0323
Pacific Data Products	619-597-3444
Pacific Rim Systems Inc.	510-782-1013
Panasonic	800-222-0584
Plextor	800-866-3935
PLI	612-830-9069
Procom Technology	800-800-8600
Proteon	508-898-3100
Pure Data	800-661-8210
Quantum hard drive	800-826-8022
Quark	303-894-8899/303-894-8822
Radius	408-541-5700
Samsung	800-446-0262
Seagate	408-438-8222
Seiko	408-922-1917
Sharp	800-BESHARP
Shiva	617-252-6400
Simple Technology	800-847-4675
SMC	800-992-4762
Sony	408-894-0225
Storage Dimensions	408-894-1325
Summagraphics	800-729-7866
Supra	503-967-2490 (WIN/DOS)
Supra	503-967-2492 (MAC)
Symantec	800-441-7234
T/Maker	415-962-0195
Teac	213-726-0303
Texas Instruments	800-848-3927
Toshiba	800-999-4273
Trantor Adeptec	408-934-7274
Tripplite	312-329-1601
Umax	800-562-0311
US Robotics	708-982-5151
Wacom	800-922-6635
Western Digital	800-275-4932
Western Digital hard drive 24-hr FAX back service	714-932-4300
WordPerfect	800-321-3389
Wyse	408-435-2770
Xircom	805-376-9200
Zenith	708-808-4300